Analysis of Human-Robot Interaction for Urban Search and Rescue

Holly A. Yanco, Michael Baker, Robert Casey, Brenden Keyes, Philip Thoren
University of Massachusetts Lowell
One University Ave, Olsen Hall
Lowell, MA, USA
{holly, mbaker, rcasey, bkeyes, pthoren}@cs.uml.edu

Jill L. Drury
The MITRE Corporation
Mail Stop K320
202 Burlington Road
Bedford, MA, USA
jldrury@mitre.org

Douglas Few, Curtis Nielsen and David Bruemmer
Idaho National Laboratories
2525 N. Fremont Ave.
Idaho Falls, ID, USA
{david.bruemmer, douglas.few, curtis.nielsen}@inl.gov

Abstract—**This paper describes two robot systems designed for urban search and rescue (USAR). Usability tests were conducted to compare the two interfaces developed for human-robot interaction (HRI) in this domain, one of which emphasized three-dimensional mapping while the other design emphasized the video feed. We found that participants desired a combination of the interface design approaches. Additionally, participants desired a combination of the interface design approaches, however, we also observed that sometimes the preferences of the participants did not correlate with improved performance. The paper concludes with recommendations from participants for a new interface to be used for urban search and rescue.**

I. INTRODUCTION

Over the past several years, two different interfaces for human-robot interaction (HRI) have been built upon a similar robot base with similar autonomy capabilities: one at the Idaho National Laboratories (INL) and the other at the University of Massachusetts Lowell (UML). The INL interface includes a three-dimensional representation of the system's map and the robot's placement within that map, while the UML system has only a two-dimensional map in its video-centric design. To learn which interface design elements are most useful in different situations, we conducted usability studies of the two robot systems at the urban search and rescue (USAR) test arena at the National Institute of Standards and Technology (NIST) in Gaithersburg, MD.

This paper presents the two robot systems and their interface designs, the experiment and analysis methodologies, the results of the experiments and strategies for designing more effective USAR interfaces. Beyond urban search and rescue, we feel the results will be relevant to the design of remote robot interfaces intended for search or monitoring tasks.

II. ROBOT SYSTEMS

This section describes the robot hardware, autonomy modes and the interfaces for the INL and UML systems.

A. Idaho National Laboratories

The INL control architecture is the product of an iterative development cycle where behaviors have been evaluated in the hands of users [2], modified, and tested again. The INL has developed a behavior architecture that can port to a variety of robot geometries and sensor suites. This architecture, called the Robot Intelligence Kernel, is being used by several HRI research teams throughout the community. The experiments discussed in this paper utilized the iRobot ATRV-Mini (shown in Figure 1), which has laser and sonar range finding, wheel encoding, and streaming video.

Using a technique described in Pacis et al. [10], a guarded motion behavior permits the robot to take initiative to avoid collisions. In response to laser and sonar range sensing of nearby obstacles, the robot scales down its speed using an event horizon calculation, which measures the maximum speed the robot can safely travel in order to come to a stop approximately two inches from the obstacle. By scaling down the speed in many small increments, it is possible to insure that, regardless of the commanded translational or rotational velocity, guarded motion will stop the robot at the same distance from an obstacle. This approach provides predictability and ensures minimal interference with the operator's control of the vehicle. If the robot is being driven near an obstacle rather than directly towards it, guarded motion will not stop the robot, but may slow its speed according to the event horizon calculation.

Various modes of operation are available, affording the robot different types of behavior and levels of autonomy. These modes include Teleoperation where the robot takes no initiative, Safe Teleoperation where the robot takes initiative to protect itself and the local environment, Standard Shared Mode where the robot navigates based upon understanding of the environment, yet yields to human joystick input, and Collaborative Tasking Mode where the robot autonomously

Figure 1: The INL robot: an iRobot ATRV-Mini

creates an action plan based on the human mission-level input (e.g. go to a point selected within the map, return to start, go to an entity).

Control of the INL system is actuated through the use of an augmented virtuality [3] 3D control interface that combines the map, robot pose, video, and camera orientation into a single perspective of the environment [8, 9]. From the 3D control interface, the operator has the ability to place various icons representing objects or places of interest in the environment (e.g. start, victim, or custom label). Once an icon has been placed in the environment, the operator may enter into a Collaborative Task by right-clicking the icon which commissions the robot to autonomously navigate to the location of interest. The other autonomy modes of the robot are enacted through the menu on the right side of the interface.

As the robot travels through the remote environment it builds a map of the area. Through continuous evaluation of sensor data, the robot attempts to keep track of its position with respect to its map. As shown in Figure 2, the robot is represented as the red vehicle in the 3D control interface. The virtual robot is sized proportionally to demonstrate how it fits into its environment. Red triangles will appear if the robot is blocked and unable to go in a particular direction. The user has the ability to select the perspective through which the virtual environment is used by choosing the Close, Elevated or Far button. The Default View returns the perspective to the original robot-centered perspective. The blue extruded columns are a representation of the robot's map. The map will grow as the robot travels through the environment.

B. UMass Lowell

UMass Lowell's robot platform is an iRobot ATRV-Jr research robot. This robot came equipped with a SICK laser rangefinder, positional sensors and a ring of 26 sonars. We have added front and rear pan-tilt-zoom cameras, a forward-looking infrared (FLIR) camera, a carbon dioxide (CO_2) sensor, and a lighting system (see Figure 3).

Figure 2: The INL USAR interface

The robot uses autonomy modes similar to INL's; in fact, the basis for the current mode system is INL's system. Teleoperation, safe, goal (a modified shared mode) and escape modes are available.

In the current version of the interface (see [1] for a description of the earlier system), there are two video panels, one for each of the two cameras on the robot (see Figure 4). The main video panel is the larger of the two and is where the user will focus while driving the robot. The second video panel is smaller, is placed at the top-right of the main video window, and is mirrored to simulate a rear view mirror in a car. By default, the front camera is in the main video panel, while the rear camera is displayed in the smaller rear view mirror video panel.

The robot operator has the ability to switch camera views, in what we call the Automatic Direction Reversal (ADR) mode. In ADR mode, the rear camera is displayed on the main video panel, and the front camera is on the smaller panel. All the driving commands and the range panel (described below) are reversed. Pressing forward on the joystick in this case will cause the robot to back up, but to the user, the robot will be moving "forward" (i.e., the direction that their current camera is looking). This essentially eliminates the front/back of the robot, and cuts down on rear hits, because the user is now very rarely "backing up."

The main video panel displays text identifying which camera is currently being displayed in it and the current zoom level of the camera (1x - 16x). The interface has an option for showing crosshairs, indicating the current pan and tilt of the camera.

Information from the sonar sensors and the laser rangefinder is displayed in the range data panel located directly under the main video panel. When nothing is near the robot, the color of the box is the same gray as the


Figure 3: The UML robot: an iRobot ATRV-JR

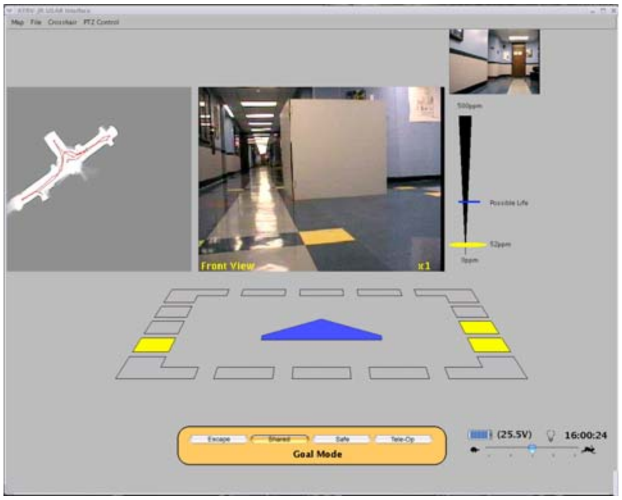

Figure 4: The UML USAR interface

background of the interface, to indicate nothing is there. As the robot approaches an obstacle at a 1 ft distance, the box will turn to yellow, and then red when the robot is very close (less than .5 ft). The ring is drawn in a perspective view, which makes it look like a trapezoid. This perspective view was designed to give the user the sensation that they are sitting directly behind the robot. If the user pans the camera left or right, this ring will rotate opposite the direction of the pan. If, for instance, the front left corner turns red, the user can pan the camera left to see the obstacle, the ring will then rotate right, so that the red box will line up with the video showing the obstacle sensed by the range sensors. The blue triangle, in the middle of the range data panel, indicates the true front of the robot. The system aims to make the robot's front and back be mirror images, so ADR mode will work the same with both; however, the SICK laser, CO_2 sensor, and FLIR camera only point towards the front of the robot, so this blue arrow helps the user to distinguish front and back if needed.

The mode indicator panel displays the current mode that the robot is in. The CO_2 indicator, located to the right of the main video, displays the current ambient CO_2 levels in the area. As the levels rise, the yellow marker will move up. If it is above the blue line, then there is possible life in the area. The bottom right of the interface has the status panel. This consists of the battery level, current time, whether the lights are on or off, and the maximum speed level of the robot.

The robot is controlled via joystick. In order for the robot to move, the operator must press the trigger, and then give it a direction. If the user presses the joystick forward, the robot will move forward, left for left, etc. On top of the joystick is a hat sensor with can read eight compass directions. This sensor is used to pan and tilt the camera. By default, pressing up on this sensor will cause the camera to tilt up, likewise pressing left will pan the camera left. An option in the interface makes it so that pressing up will cause the camera to tilt down; some people, especially pilots, like this option. The joystick also contains buttons to home the cameras, perform zoom functions, and toggle the brake. It also has a button to toggle Automatic Direction Reversal mode. Finally, a scrollable wheel to set the maximum speed of the system is also located on the joystick.

III. METHODOLOGY

A. Experimental Set-Up

Because we wished to see differences in preferences and performance with the UML interface and the INL interface, we designed a within-subjects experiment with the independent variable being interface type. Eight people (7 men, 1 woman) ranging in age from 25 to 60 with search and rescue experience agreed to participate.

We asked participants to fill out a pre-experiment questionnaire so we could understand their relevant experience prior to training them on how to control one of the robots. We allowed participants time to practice using the robot in a location outside the test arena and not within their line of sight so they could become comfortable with remotely moving the robot and the camera(s) as well as the different autonomy modes. Subsequently, we moved the robot to the arena and asked them to maneuver through the area to find victims. We allowed 25 minutes to find as many victims as possible, followed by a 5-minute task aimed primarily at ascertaining situation awareness (SA). After that, we took a short break during which an experimenter asked several Likert scale questions. Finally, we repeated these steps using a different robot, ending with a final short questionnaire and debriefing. The entire procedure took approximately 2 1/2 hours.

The specific tasking given to the participants during their 25-minute runs was to "fully explore this approximately 2000 foot space and find any victims that may be there, keeping in mind that, if this was a real USAR situation, you'd need to be able to direct people to where the victims were located." Additionally, we asked participants to "think aloud" [4] during the task. After this initial run, participants were asked to maneuver the robot back to a previously seen point, or maneuver as close as they could get to it in five minutes. Participants were not informed ahead of time that they would need to remember how to get back to any particular point.

We counterbalanced the experiment in two ways to avoid confounders. Five of the eight participants started with the UMass Lowell system and the other three participants began with the INL system. (Due to battery considerations, a robot that went first at the start of the day had to alternate with the other system for the remainder of that day. UML started first in testing on days one (2 subjects) and three (3 subjects). INL started first on day two (3 subjects).) Additionally, two different starting positions were identified in the arena so that knowledge of the arena gained from using the first interface would not transfer to the use of the second interface; starting points were split changed between users. The two counterbalancing techniques led to four different combinations of initial arena entrance and initial interface.

The tests were conducted in the Reference Test Arenas for Autonomous Mobile Robots developed by the National Institute of Standards and Technology (NIST) [5, 6]. During these tests, the arena consisted of a maze of wooden partitions and stacked cardboard boxes. The first half of the arena had wider corridors than the second half.

B. Analysis Methods

Analysis consisted of two main thrusts: understanding how well participants performed with each of the two interfaces, and interpreting their comments on post-run questionnaires.

Performance measures are implicit measures of the quality of the user interaction provided to users. Under ordinary circumstances, users who were given usable interfaces could be expected to perform better at their tasks than those who were given poor interfaces. Accordingly, we analyzed the percentage of the arena explored, the number of times the participants bumped the robot against obstacles, and the number of victims found.

After each run, participants were asked to name the features that they found "most useful" and "least useful." We inferred that the "useful" features were considered by participants to be positive aspects of the interface and the "least useful" features were, at least in some sense, negative. After reviewing all of the comments from the post-run questionnaires, we determined that they fell into five categories: video, mapping, other sensors, input devices, and autonomy modes. Results are provided in the next section.

IV. RESULTS AND DISCUSSION

A. Performance Measures

1) Area Coverage: We hypothesized that the three-dimensional mapping system on INL's interface would provide users with an easier exploration phase. Table I gives the results of arena coverage for each participant with each of the robot systems. There is a significant difference ($p<.022$, using a two-tailed paired t-test with dof=7) between the amount of area covered by the INL robot and the amount covered by the UML robot, seeming to confirm our hypothesis.

One possible confounding variable for this difference is the size of the two robots. The ATRV-Mini (INL's robot) is smaller than the ATRV-Junior (UML's robot) and thus could fit in smaller areas. However, the first half of the arena, which was the primary area of coverage, had the widest areas, fitting both robots comfortably.

TABLE I

COMPARISON OF THE PERCENTAGE OF THE ARENA COVERED FOR TWO INTERFACES

	% Area Covered	
Participant	INL	UML
1	8.7	12.6
2	37.9	25.2
3	34.8	34.8
4	37.9	19.7
5	30.3	27.3
6	33.3	22.7
7	53.0	31.8
8	30.3	19.7
Average	33.3 (7.8)	24.2 (5.8)

2) Number of Bumps: One implicit measure of situation awareness is the number of times that the robot bumps into something in the environment. However, there were several confounding issues in this measure. First, the INL robot experienced a sensor failure in its right rear sensors during the testing. Second, the INL robot has a similar length and width, meaning that it can turn in place without hitting obstacles; the UML robot is longer than it is wide, creating the possibility of hitting obstacles on the sides of the robot. Finally, subjects were instructed not to use the teleoperation mode (no sensor mediation) on the INL robot, while they were allowed to use it on the UML robot.

Despite these confounding factors, we found no significant difference in the number of hits that occurred on the front of the robot (INL average: 4.0 (3.7); UML average: 4.9 (5.1); $p=.77$). Both robots are equipped with similar cameras on the front and both interfaces present some sort of ranging data to the user. As such, the awareness level of obstacles in front of the robot seems to be similar between systems.

When hits occurring in the back right of the robot were eliminated from both counts, we did find a significant difference in the number of hits (INL average: 2.5 (1.6); UML average: .75 (1.2); $p<.037$). The UML robot has a camera on the rear of the robot, adding additional sensing capability that the INL robot does not have. While both robot systems present ranging information from the back of the robot on the interface, the addition of a rear camera appears to improve awareness of obstacles behind the robot.

The systems also had a significant difference in the number of hits on the side of the robot (INL average: 0 (0); UML average: 0.5 (0.5); $p<.033$). As the two robots had equivalent ranging data on their sides, the difference in hits appears to come solely from the robot's size and geometry.

3) Victims Found: We had hypothesized that the emphasis on the video window and other sensor displays such as the FLIR and CO_2 sensor of the UML interface would allow for users to find more victims in the arena. However, this hypothesis was not borne out by the data because there was an insignificant difference ($p=.35$) in the number of victims found. Using the INL system, participants found an average of .63 (.74) victims. With the UML system, participants found an average of 1.0 (1.1) victims.

In general, victim placement in the arena was sparse and the victims that were in the arena were well hidden. Using the number of victims found as an awareness measure might have been improved by a larger number of victims, with some easier to find than others.

B. User Preferences

1) Likert scale: At the end of each run, users were asked to rank the ease of use of each interface, with 1 being extremely difficult to use and 5 being very easy to use. In this subjective evaluation, operators found the INL interface more difficult to use: 2.6 for INL vs 3.6 for UML ($p = .0185$).

Users were also asked to rank how the controls helped or hindered them in performing their task, with 1 being "hindered me" and 5 being "helped me tremendously." Operators felt that the UML controls helped them more: 4.0 for UML and 3.2 for INL ($p=.0547$).

2) Interface Features: Users were also asked what features on the robots helped them and which features did not. We performed an analysis of these positive and negative statements, clustering them into the following groups: video, mapping, sensors, input devices and autonomy. The statements revealed insights into the features of the systems that the users felt were most important.

In the mapping category, there were a total of 10 positive mapping comments and one negative for the INL system and 2 negative mapping comments overall for the UML system. We believe that the number of comments shows that the participants recognized the emphasis on mapping within the INL interface and shows that the three-dimensional maps were preferred to the two-dimensional map of the UML interface. Furthermore, the preference of the INL mapping display and the improved average percentage of the environment covered by the INL robot suggests that the user preferences were in-line with requirements for improved performance. Interestingly, two of the positive comments for INL identified the ability to have both a three-dimensional and two-dimensional map. Subjects also liked the waypoint marking capabilities of the INL interface.

There were a similar number of comments made on video about the two systems (13 for UML and 16 for INL). This seems to suggest that video is very important in this task, and most subjects were focused on having the best video possible. There were more positive comments for UML (10 positive and 3 negative) and more negative comments for INL (3 positive and 13 negative). The INL video window moved when the camera was panned or tilted; the robot stayed in a fixed position within the map while the video view moved around the robot. This video movement caused occlusion and distortion of the video when the camera was panned and tilted, making it difficult to use the window to identify victims or places in the environment. It is of interest that despite the feelings by many participants about how the video should be presented, there was no significant difference ($p=.35$) in performance with respect to the number of victims found. This disconnect between preference and performance suggests that more work is required to understand what presentation of the video will actually improve the operator's ability to search an environment.

Interestingly, most of the positive video comments for UML did not address a fixed position window (only 1 comment). Four users commented that they liked the ability to home the camera (INL had two positive comments about this feature as well). Three users commented that they liked having two cameras.

All comments on input devices were negative for both robots, suggesting that people just expect that things will work well for input devices and will complain only if they aren't working. There were a similar number of positive comments for autonomy, suggesting that users may have noticed when the robot had behaviors that helped. It is possible that the users didn't know what to expect with a robot and thus were just happy with the exhibited behaviors and accepted things that they may not have liked.

We saw many more comments on UML's sensors (non-video), which identifies the emphasis on adding sensors on the UML system. INL had two negative comments for not having lighting available on their robot. UML had 10 positive comments (1 each for lights, FLIR and CO_2, 4 for the laser ranging display and 3 for the sonar ring display) and 3 negative comments (2 for the sonar ring display blocks not being definitive and 1 for the FLIR camera).

Our analysis suggests that there are a few categories of great importance to operators: video, labeling of maps, ability to change perspective between 3D and 2D maps, additional sensors, and autonomy. In fact, in their suggested ideal interface, operators focus on these categories.

C. Designing the Ideal Interface

After using both interfaces, users were asked which features they would include if they could combine features of the two interfaces to make one that works better for them. Every user had his or her own opinion, as follows:

- Subject 1 wanted to combine map features (breadcrumbs on the UML interface and labeling available on the INL interface).
- Subject 2 wanted to keep both types of map view (3D INL view, 2D UML view), have lights and add other camera views (although this user also remarked that he didn't use UML's rear view camera much).
- Subject 3 wanted to add the ability to mark waypoints to the UML system.
- Subject 4 liked the blue blocks on INL (3D map walls), the crosshairs on UML (pan and tilt indicators on the video), the stationary camera window on the UML interface, marking entities and going to waypoints on the INL interface, the breadcrumbs in the UML map, and the bigger camera view that the UML interface had.
- Subject 5 liked the video set up on UML and preferred the features on the UML interface. He would not combine any features.
- Subject 6 wanted a fixed camera window (like UML), a 2D map in the left hand corner of the 3D interface, the ability to mark waypoints on the map, roll and pitch indicators, and lights on the robot.
- Subject 7 wanted to take UML as a baseline interface, but wanted a miniaturized blue block map (3D map) instead of the 2D map, since it provided more scale information.
- Subject 8 wanted to start with the UML interface, with the waypoint marking feature and shared mode capability of the INL system.

When asked to design their ideal interface, most subjects commented on the maps, preferring the 3D map view to the 2D view; the 3D map view provides more information about the robot's orientation with respect to the world. Features of the two maps could be combined, either with a camera view that could swing between 3D and 2D or by putting both types of maps on the screen. However, operators did comment that they did not like the way that the current implementation of the blue blocks obscured the video window when it was tilted down or panned over a wall.

Most subjects also expressed a desire to have an awareness of where they had been, with the ability to make annotations to the map. They wanted to have the "breadcrumbs" present on the UML interface, which showed the path that the robot had taken through the arena. This feature was available on the INL interface, but not turned on for the experiments. Subjects also wanted to be able to mark waypoints on the map, which was a feature in the INL system.

The subjects did not like the moving video window present on the INL interface, preferring a fixed camera window instead. We believe that in a USAR task, a fixed window of constant size allows for the operator to more effectively judge the current situation. While this hypothesis seems to be borne out by the comments discussed above, it was not verified by measures such as number of victims found and number of hits in the front of the robot, both of which were not statistically different between the two systems.

Interestingly, when designing their interface, no subjects commented on the additional sensors for finding victims that were present on the UML system: the FLIR camera and the CO_2 sensor. It seemed that their focus fell on being able to understand where they were in the environment, where they had been, and what they could see in the video.

V. CONCLUSIONS

Eight trained USAR personnel tested two robot systems. The purpose of the experiment was to understand how the robot systems affected the operator's ability to perform a search task in an unknown environment. The two robot systems utilized different physical robots and control algorithms as well as different interfaces and sensor suites.

From the experiment, it was observed that the camera information was particularly important to the operators because many of their likes and dislikes concerned the presentation of the video information. However, it is of note that despite the subjective preferences of the operators, there was not a significant difference in the number of victims found. Furthermore, it was observed that the search task was largely unsuccessful as, on average, less than one of four victims was found. Improvement of technology and evaluation techniques will be necessary to answer the question of what improves performance in search tasks.

The occlusion of video by other sets of information may have influenced the operator's ability to adequately search the environment, as it was more difficult for the operator to see the entire visual scene. Another possibility is that the navigational requirement of the task took sufficient effort from the participant that it negatively impacted the operator's ability to search the environment. Even though there were various levels of autonomy available to facilitate the navigation of the robot, participants often expressed confusion about where the robot had been and what they had seen previously. To improve the usefulness of robot systems in search and detection tasks in general, it will be important to reduce the operator's responsibility to perform both the navigation and search aspects of the task.

VI. FUTURE WORK

There are two efforts currently under investigation that are the result of the experiments described in this paper. The first effort is a method that will enable operators to focus on the search aspect of the task by minimizing his or her responsibility in the navigating through the remote environment. Although previous work has sought to reduce the human's navigational responsibility by improving the robot's navigational autonomy, it left the navigation and exploration tasks as separate processes that both required a level of operator attention. The new approach currently being investigated integrates the navigational task into the search task by providing a "navigate-by-camera" mode. In this mode, the operator directs the camera to points of interest and the robot maneuvers to them while avoiding obstacles and keeping the camera focused on the specified point. This mode should help the operator by allowing them to focus on where the camera is pointing and not how to get the robot from place to place.

The second effort being investigated is to help the operator understand where they have and have not searched within the remote environment. To do this, we will continue the use of labels and icons, but make them more customizable so that they can include user-defined images to represent places of interest. Additionally, even though a breadcrumb trail was useful to indicate where the robot had been, it did not illustrate in three-dimensional space where the operators have looked. To increase this knowledge we are investigating the use of a representation that presents information about where the camera was pointing as the robot was moved through the environment. This should enable operators to quickly recognize what parts of the environment have been "seen" by the robot and continue on to unseen areas. Finally, to help the operator remember the environment better, we are investigating new ways to transition between ego-and exo-centric perspectives of the environment such that the transition is quick and intuitive and supports a "quick-glance" at the robot's location within a larger environment.

We anticipate that these approaches will improve the usefulness of remote robots in urban search and rescue tasks as well as other remote robot tasks that require the use of video information in conjunction with navigational information.

ACKNOWLEDGEMENTS

This work is sponsored in part by the National Science Foundation (IIS-0415224, IIS-0308186), the National Institute of Standards and Technology (70NANB3H1116), and the Idaho National Laboratory's Intelligent Systems Initiative.

REFERENCES

[1] M. Baker, R. Casey, B. Keyes and H. A. Yanco. "Improved interfaces for human-robot interaction in urban search and rescue." In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, The Hague, The Netherlands, October 2004.

[2] D.J. Bruemmer, D.A. Few, R.L. Boring, J.L. Marble, M.C. Walton, and C. W. Nielsen. "Shared Understanding for Collaborative Control." In *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*. Volume 35 Number 4, pp 494-504, July 2005

[3] D. Drascic and P. Milgram. "Perceptual issues in augmented reality." In *Proceedings of SPIE Vol. 2653: Stereoscopic Displays and Virtual Reality Systems III*, San Jose, CA, 1996.

[4] K. A. Ericsson and H. A. Simon. "Verbal reports as data." *Psychological Review,* Vol. 87, pp. 215 – 251, 1980.

[5] A. Jacoff, E. Messina, and J. Evans. "A reference test course for autonomous mobile robots." In *Proceedings of the SPIE-AeroSense Conference*, Orlando, FL, April 2001.

[6] A. Jacoff, E. Messina, and J. Evans. "A standard test course for urban search and rescue robots." In *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, August 2000.

[7] C. W. Nielsen, B. Ricks, M. A. Goodrich, D. J. Bruemmer, D. A. Few, and M. C. Walton. "Snapshots for semantic maps." In *Proceedings of the 2004 IEEE Conference on Systems, Man, and Cybernetics*, The Hague, The Netherlands, 2004.

[8] C. W. Nielsen and M. A. Goodrich. "Comparing the usefulness of video and map information in navigation tasks." In *Proceedings of the Human Robot Interaction Conference*. Salt Lake City, UT, 2006.

[9] C. W. Nielsen, M. A. Goodrich, and R. J. Rupper. "Towards facilitating the use of a pan-tilt camera on a mobile robot." In *Proceedings of the 14th IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN)*, Nashville, TN, 2005.

[10] E.B. Pacis, H.R. Everett, N. Farrington, and D. J. Bruemmer. "Enhancing Functionality and Autonomy in Man-Portable Robots." In P*roceedings of the SPIE Defense and Security Symposium 2004*. 13 -15 April, 2004.